

**IN THE UNITED STATES
PATENT AND TRADEMARK OFFICE**

TITLE:

Apparatus and Method for Managing the Performance
of an Electronic Device

INVENTORS:

Howard Marshall Monroe

Thomas Alan Mee

ASSIGNEE:

United Services Automobile Association
(USAA)

[0001] This utility application claims priority on a United States Provisional Application entitled "Multiple Exponential Adaptive Triage," Serial No. 60/436,169, having a filing date of December 23, 2002.

FIELD OF THE INVENTION

5

[0002] The present invention relates generally to resource management and, more particularly, to a system and method of predicting and managing the performance of one or more electronic devices.

BACKGROUND OF THE INVENTION

10 [0003] Managing a computer system which includes a plurality of devices such as networks or servers, is of special interest to data processing or information technology personnel. Such computer systems typically include a plurality of diverse devices including memory, disks, local area network (LAN) adaptors and central processing units (CPUs) which interact in various ways to facilitate data processing applications.

15 [0004] As systems become larger and more complex, interactions between electronic devices become harder to define, model, and/or predict. Such systems may suffer from inefficiencies or "bottlenecks" that slow or even stop the system.

[0005] Often, the performance of a computer system or network is less than it could be because of one or more components having an inappropriate load applied thereto. Thus, it is desirable to know

what changes to the system would be required in order to improve capacity of each electronic device. Further, such changes would allow the manipulation of a preset number of electronic devices instead of the system as a whole. To accomplish this, there remains a need for a system and method capable of collecting and analyzing performance data such that it may be utilized to predict future
5 performance of individual electronic devices.

SUMMARY OF THE INVENTION

[0006] Accordingly, the present invention provides a performance management system and method capable of generating a plurality of forecasts for one or more electronic devices. The forecasts are
10 generated from stored performance data and analyzed to determine which devices are likely to experience performance degradation within a predetermined period of time. A single forecast is chosen for further analysis such that computer modeling may be performed upon the performance data to enable the user to predict when device performance will begin to degrade. In one embodiment, graphical displays are created for devices forecasted to perform at an undesirable level
15 within a predetermined period of time.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] A more complete appreciation of the invention and many of the attendant advantages thereof will be readily obtained as the same becomes better understood by reference to the following
20 detailed description when considered in connection with the accompanying drawings, wherein:

[0008] Figure 1 is a component diagram of one embodiment of the present invention.

[0009] Figures 2, 3, 5 and 6 are process flow diagrams illustrating various embodiments of the present invention.

[0010] Figure 4 is a graphical display illustrating the data modeling and analysis process of one
5 embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

[0011] The present invention is herein described as a method of managing one or more electronic devices and as a computer system for managing one or more electronic devices.

10 [0012] Referring to Figure 1, the computer system (10) of the present invention is capable of receiving and analyzing data from any number of electronic devices (12). In one embodiment, data describing the performance of such devices is collated and processed by an intermediate processing unit (14) prior to the storage of the data upon a storage device (16).

[0013] In another embodiment, performance data is fed through a computer network for storage
15 upon the storage device. In this embodiment, the data is collated and processed by a central processing unit (14) coupled to each electronic device, as well as the computer network.

[0014] In one embodiment, the processing unit (14) of the present invention is equipped with a graphic display interface (20) capable of providing graphical displays of analyzed performance data, as discussed further below.

[0015] In one embodiment, the results of analysis performed by the processing unit may be sent either to a printer (22) for the creation of hard copy reports, or electronically to one or more analysis personnel (24). In one embodiment, such analyzed information may be transmitted through one or more computer networks (18). Further, the reporting capabilities of the present invention allow this system to provide the analyst with analysis summaries. This feature of the present invention provides the analyst with an overview of one or more of the electronic devices at issue, in order to allow the analyst to make an informed decision regarding which devices require attention.

[0016] For the purposes of illustration only, in one embodiment, the present invention may be utilized to determine when to upgrade an Intel® server, such as a Compaq® Pentium II® having a quad processor running at 333MHz. A performance data collection tool residing on the server, Best 1 for example, is capable of capturing performance data every few seconds. Data is then sent to an IBM R/S 6000 midrange server, via a local area network (LAN), where it is collected and processed. A batch job is then run, using SAS® Proc Reg for example, which appends the data into a database stored on a storage area network (SAN). The data may then be gathered from the SAN and analyzed according to the present invention using, for example, a Dell® Desktop computer having a Pentium IV® processor operating at 1.7GHz. Capacity charts may be displayed on the computer's monitor, sent to a printer, and/or stored electronically on the SAN.

[0017] Referring to Figures 2 and 3, the present invention is capable of collecting data from a host of electronic devices in order to determine potential performance degradation over time. Referring to box (26), the present invention extracts and analyzes performance data held upon a storage device (16). In one embodiment, performance data is collected and analyzed by an intermediate processing
5 unit (14I), described above. During processing, the data may be formatted such that it may be analyzed by any number of known statistical analysis systems. In one embodiment, performance data is analyzed by Statistical Analysis System® (SAS) software capable of applying a host of statistical procedures and data management tools.

[0018] Such statistical analysis systems are utilized by the present invention to generate a plurality
10 of forecasts relating to the performance of one or more electronic devices. In one embodiment, the present invention is utilized to analyze one or more servers such that a plurality of forecasts may be generated for each one, as illustrated by box (28).

[0019] The present invention may utilize any number of known statistical methods, in order to generate a plurality of forecasts for each device. In one embodiment, the system (10) of the present
15 invention generates graphical displays of each forecast for review by the user, as illustrated by box (32). In one embodiment, such displays may be used to provide the user with an overview of a device's capacity as well as fluctuations over any given period of time.

[0020] The processing unit (14) of the computer system (10) of the present invention selects a single forecast from the plurality of forecasts generated, as illustrated by box (30). In one embodiment, this
20 is accomplished via the assignment of selection parameters by the user or analyst, as illustrated by

box (34). These parameters may consist of a threshold value relating to the device data being analyzed or a predetermined time period. For example, the graphical display of Figure 4, illustrates “device capacity” on the Y-axis, and “time” on the X-axis. In this example, a capacity threshold of 11 SPECint95, an industry standard performance benchmark, has been selected in order to enable the system to readily ascertain capacity readings above this threshold value. By setting such a threshold value, the user may instruct the system to single out forecasts showing capacity readings above or below a preset figure. In one embodiment, the threshold value is determined to be the point at which the capacity of the device or devices in question begin to degrade. In another embodiment, a threshold capacity of 70% measured in relation to the maximum capacity of the device is utilized.

[0021] Multiple selection parameters may be assigned to enable the system to single out individual forecasts having particular attributes. For example, if the user assigns a capacity threshold of 11 SPECint95 and a planning horizon of January, 2003, the graphical display of Figure 4 would be flagged by the system such that further review may be conducted by an analyst. Specifically, the selected capacity of 11 SPECint95 is forecasted to exceed prior to the planning horizon date of January, 2003. Thus, further review of the device at issue is warranted, given the selection parameters.

[0022] Figure 4 graphically illustrates the actual capacity data (36), the selected capacity threshold (38), the total capacity (40), and the forecasted performance of a device (42) in light of collected performance data. In one embodiment, the planning horizon is chosen based on the estimated time it would take to repair and/or upgrade the device at issue. Further, in one embodiment, the single

forecast selected by the system is the most conservative forecast relating to the device upon which statistical analysis has been performed. Thus, the present invention allows the user to save time and effort by reviewing only those forecasts indicating an urgent need for review, as illustrated by boxes (44) and (46) of Figures 5 and 6. This process may then be repeated for additional individual
5 devices or an aggregate of devices, as illustrated by box (48).

[0023] Data modeling for the purpose of generating forecasts, as illustrated by box (52), is well known in the art. For the purposes of the present invention, any reasonable linear or non-linear statistical model (such as a polynomial or an exponential model) may be utilized to forecast when the capacity threshold of an electronic device will be exceeded. In one embodiment, a log-linear
10 model is utilized because it provides conservative forecasts for many electronic devices. In many cases, the resource consumption growth of an electronic device will behave in a log-linear fashion. Thus, for many electronic devices, the log-linear model provides the most accurate forecasting results.

[0024] This model takes the form of $Y = Ac^{Bx}$, where Y is the actual capacity consumed, X is the
15 time or date of consumption, c is the rate of exponential change, and A and B are model parameters to be estimated from the data. One may assume any value for c, or, in one embodiment, estimate the value of c from actual data. However, in this example, c is taken to be 2.71828.

[0025] In one embodiment, the forecasting procedure utilized by the present invention for this example model begins by converting the exponential model to natural logarithms, i.e., $\ln Y = \ln A + Bx$, and then utilizing linear regression to determine estimates of $\ln A$ and B. In order to
20

determine an estimate for Y, the present invention calculates $e^{\text{Ln } A + Bx}$ by replacing Ln A and B with their respective estimates. The use of linear regression is well known in the art and is available in many statistical analysis computer packages.

[0026] For the data shown in the example of Figure 4, the estimate of Ln A equals -59.46628, and
5 the estimate of B equals 0.00396. For example, if the date x is taken to be August 1, 2001, the estimated capacity consumed would equal 1.97. Thus, the value of x utilized by the example of Figure 4 is 15,188, which is the number of days from January 1, 1960, to August 1, 2001. As mentioned above, models different from the exponential method described above would require different methods for estimating the model parameters. It being understood that the present
10 invention is capable of utilizing any of such statistical models during the analysis process.

[0027] Forecasting techniques may be applied either to peak, average, ninety-fifth percentile, minimum, or any other chosen statistic in order to provide a statistical confidence interval for each forecast. Further, the present invention may be applied to software applications executed by one or more electronic devices. In one embodiment, forecasts may be sorted by date in order to inform the
15 user or administrator which device(s) require the most immediate attention. Further, each forecast may be reviewed and validated by one or more analysts familiar with the device at issue, as illustrated by box (50). In some cases, the analyst agrees with the forecasted data and begins an upgrade and/or adjustment of the device to prevent performance degradation. In other cases, forecasted data may be in error such that the analyst recommends that the forecast be disregarded, as
20 illustrated by box (49).

[0028] Subsequent to the review of the selected forecasts by the analyst, the selection parameters discussed above may be adjusted according to the modeling results, as illustrated by box (54).

Further, the statistical analysis methodology utilized by the present invention may be adjusted based upon the knowledge and experience of the analyst with respect to the device at issue. In one embodiment, additional statistical analysis as illustrated by box (56) is conducted utilizing adjusted selection parameters to create a subsequent graphical display for further review by the analyst, as illustrated by box (32). Different types of electronic devices may require different statistical analysis methodologies. For example, if the user desires to analyze a large system comprising a plurality of different types of devices, the planning horizon and threshold capacity would require adjustment.

Further, the statistical methods utilized for each type of device may require adjustment in order to produce the most conservative forecast results. For example, a first statistical analysis method may take all of the input data into account such that each data point is weighted equally. However, a second data analysis technique may weigh the most recent data points more heavily, depending on which technique is being utilized. Thus, depending on the type of statistical analysis technique employed, the forecasted results will change accordingly.

[0029] In one embodiment, the present invention allows the user to choose from a host of statistical analysis methods complete with explanation as to what device types are most suited for each of said statistical analysis methods. Further, in one embodiment, the present invention provides the analyst with guidelines regarding the time required to upgrade/adjust various types of electronic components. This feature of the present invention allows the analyst to easily choose his or her

selection parameters, thus providing the analyst or other user with enhanced efficiency and ease of use.

[0030] Further, the processing unit of the present invention is capable of determining the device type based upon the format of the data collected and stored upon the storage device, as illustrated by box 5 (58). In one embodiment, the metadata characteristics of the performance data for each electronic device is utilized by the processing unit to determine the type of device at issue. Once the device type has been ascertained, the system is capable of automatically, or through manual direction, selecting and utilizing the most appropriate statistical analysis method and/or selection parameters suitable for the device type at issue, as illustrated by box (60). The above analysis process may then 10 be repeated utilizing the methods/parameters best suited to the device(s) at issue.

[0031] It should be understood that the present invention is not relegated to the use of capacity data or any other particular number or type of device, as in the above example. On the contrary, the present invention may utilize any metric, or combination of metrics, such as intensities, in place of or in addition to system utilization data. For example, electronic device usage may fluctuate according 15 to traffic and/or maintenance patterns. The present invention allows the analyst to adjust the forecasting results to compensate for peak and off-peak workload patterns, as well as maintenance or resource allocation procedures. This information may be stored on the storage device for use in conjunction with performance data collected from the device(s) during statistical analysis. Further, such information, through the display and reporting capabilities of the present invention, may assist 20 the analyst in consolidating and/or adjusting workload or load balancing parameters.

[0032] This feature of the present invention is particularly useful when applied to a large number of electronic devices undergoing aggregate statistical analysis. Specifically, selected forecasts may be aggregated in order to present the analyst with a performance overview of analyzed devices working in combination. This allows the analyst to conduct load balancing and/or consolidate workloads over multiple devices. The reporting capabilities described above may be utilized to provide the analyst with best and/or worst case forecasts designed to provide an overall “level of confidence”. In one embodiment, this is accomplished through graphical and/or textual display of performance data.

[0033] Although the invention has been described with reference to specific embodiments, this description is not meant to be construed in a limited sense. Various modifications of the disclosed embodiments, as well as alternative embodiments of the invention will become apparent to persons skilled in the art upon the reference to the description of the invention. It is, therefore, contemplated that the appended claims will cover such modifications that fall within the scope of the invention.